NATIONAL INTEGRATED CYBERINFRASTRUCTURE SYSTEM
SANReN

SANReN's 100 Gbps Data Transfer Service: Transferring data fast!

Short paper by: The SANReN large data transfer team: Kasandra Pillay, Johann Hugo, Ajay Makan, Thokozani Khwela, Thuso Bogopa and Manqoba Shabalala

Presented by Dr Kasandra Pillay
SC24, Atlanta, Georgia

A national initiative of the Department of Science and Innovation and implemented by the CSIR

science & innovation
Department:
Science and Innovation
REPUBLIC OF SOUTH AFRICA

CSIR
Touching lives through innovation

NICIS
NATIONAL INTEGRATED CYBERINFRASTRUCTURE SYSTEM
SANReN

AGENDA (Overview of short paper):

- About NICIS
- Background of the service
- The new 100Gbps large data transfer infrastructure of the South African National Research Network (SANReN).
- 100Gbps Data Transfer Service
- Combined DTN/perfSONAR nodes
- Benchmarking our 100Gbps DTNs
- Transfers to/from CHPC
- Types of use cases
- Supercomputing participation

A national initiative of the Department of Science and Innovation and implemented by the CSIR
science & innovation
Department:
Science and Innovation
REPUBLIC OF SOUTH AFRICA
CSIR
Touching lives through innovation

Background: NICIS

- National Integrated Cyberinfrastructure System (NICIS)
- Structure
  - South African Research Network (SANReN)
  - Centre for High Performance Computing (CHPC)
  - Data Intensive Research Initiative of South African (DIRISA)
  - HCD encompasses the 3 pillars
- NICIS is a hosted programme of the DSI
- Hosted at the CSIR as a centre in NGEI Cluster, Smart Society Division

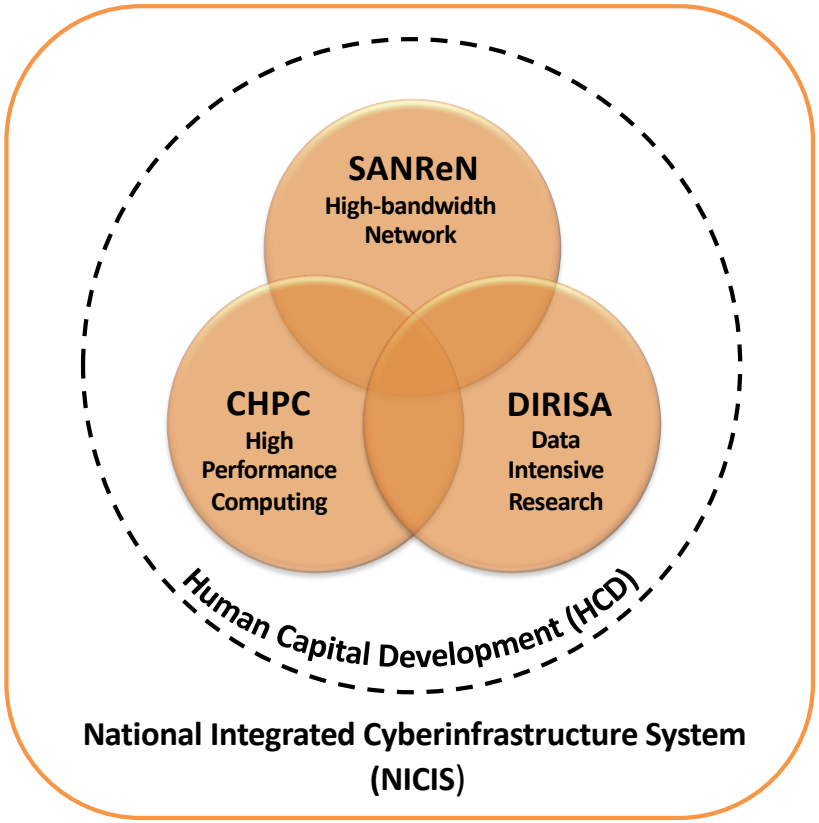

A national initiative of the Department of Science and Innovation and implemented by the CSIR

science & innovation
Department:
Science and Innovation
REPUBLIC OF SOUTH AFRICA

CSIR
Touching lives through innovation

SANReN, https://www.sanren.ac.za/ | CHPC, https://www.chpc.ac.za/ | DIRISA, https://www.dirisa.ac.za/

The South African NREN

- Structure
  - Roles and responsibilities of the de facto South African NREN are distributed between SANReN and TENET
- Functions
  - SANReN
    - Builds the Network (network reach and capacity)
    - Develops Advanced Services
  - TENET
    - Operates the SANReN Network under terms of collaboration agreement with CSIR
    - Build onto the network
    - Host Services
    - Represents South African NREN at UbuntuNet Alliance (founding member) – Regional REN for South East Africa
    - SABEN – entity for the connection of TVET colleges

A national initiative of the Department of Science and Innovation and implemented by the CSIR

science & innovation
Department:
Science and Innovation
REPUBLIC OF SOUTH AFRICA

CSIR
Touching lives through innovation

TENET, https://www.tenet.ac.za/ | SANReN, https://www.sanren.ac.za/

South African NREN: Network

South African NREN Backbone Map: Terrestrial and undersea capacity

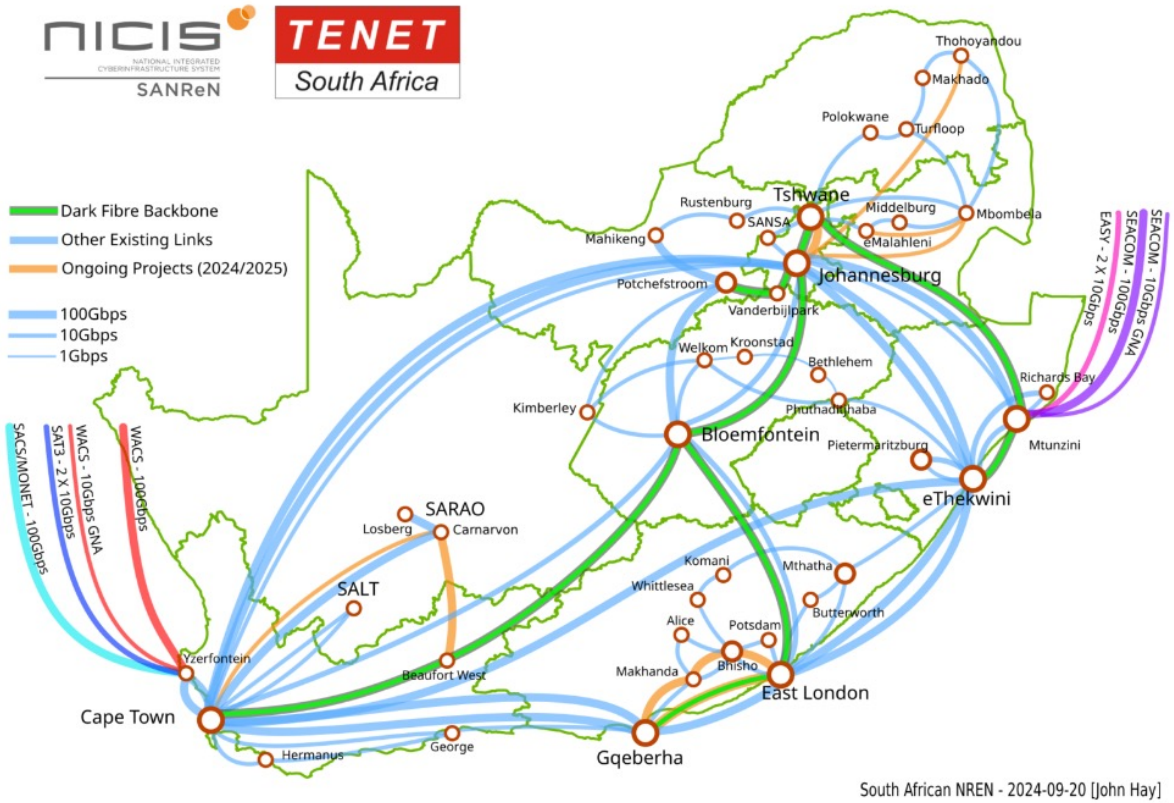

Infrastructure:

- a core national dark fibre backbone with several managed bandwidth backbone links at 100Gbps to mesh up the backbone
- backbone extensions (regional links) – typically at 10Gbps
- back-hauling from the submarine cable landing stations at Yzerfontein and Mtunzini
- capacity on undersea cables
- several metropolitan area networks

International capacity

- West African Cable System (WACS)
- South Atlantic 3 (SAT-3)
- Eastern Africa Submarine System (EASSy)
- SEACOM
- South Atlantic Cable System (SACS)

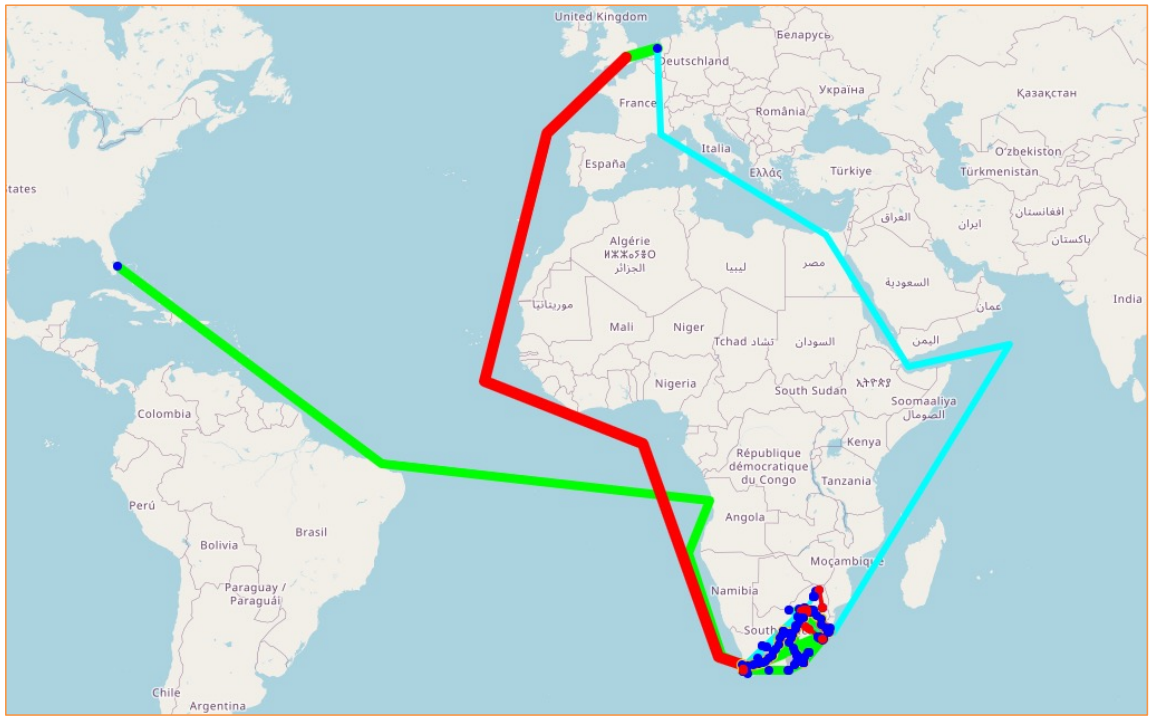

A national initiative of the Department of Science and Innovation and implemented by the CSIR

science & innovation
Department:
Science and Innovation
REPUBLIC OF SOUTH AFRICA

CSIR
Touching lives through innovation

South African NREN link status: https://linkmon.sanren.ac.za/nrenmap/

Background

- The data transfer problem – shipping disks, not leveraging the network capacity available, restricted by business networks on campus
- Challenges: Costs, Expertise, Power
- Therefore, there was a need for reduced and shared infrastructure, especially in Africa
- 2018 – Data Transfer Pilot initiated - To demonstrate data transfer improvements and save on individual costs
- Having access to high-speed networks and specialised software and hardware solutions allow us to leverage the networks capabilities for large data transfers, no more shipping hard drives, researchers should be able to access very large data sets for their research quickly and efficiently

The Solution – shared 100G perfSONAR/DTN infrastructure

NATIONAL INTEGRATED CYBERINFRASTRUCTURE SYSTEM
SANReN

Design Criteria:

- 1U device that can replace the current SANReN 10 Gbps perfSONAR server;
- capable of acting as a Globus DTN node, or a perfSONAR node, or a combined DTN/perfSONAR node;
- have dual 100 Gbps network adapters, capable of a combined 200Gbps network bandwidth;
- have Linux OS disk mirroring enabled; and
- have fast data storage, upgradeable from 10TB to 50TB (for the DTN option).

Further requirements

- The motherboard and all its components needed to support PCIe v4.0.
- Many PCIe lanes to provide enough bandwidth between the motherboard and the expansion cards – to allow for future expansion
- Data storage should use NVMe SSDs for best I/O performance.
- The selected components are very similar to those used by Netflix for its AMD-based streaming servers.

Novelty of combined DTN/perfSONAR node

- **Network perspective:** a DTN node and a perfSONAR node are located at the same networking level, and swapping physical locations will not influence performance.
- **Hardware perspective:** requirements almost identical, except for additional memory and storage requirements needed for a DTN node.
- **Budget perspective:** beneficial if these two functions can be combined on a single hardware platform – also avoiding high rack space cost.
- **Space perspective:** preferred at locations with limited rack space

The SANReN DTN/perfSONAR hardware have the following specification:

- *Server*: Supermicro AS-1114S-WN10RT;
- *CPU*: Single socket, 2nd GEN AMD EPYC 7502P Rome Processor;
- *Memory*: 128GB - 8x (16GB DDR4 3200);
- *Network*: Dual 100 Gbps Nvidia Mellanox ConnectX-6; and
- *Disks*: PCI Express 4.0 NVMe. KCD6XLUL1T92

A national initiative of the Department of Science and Innovation and implemented by the CSIR

Major constraint of a combination node

- A **major constraint** to combine perfSONAR with Globus is that **each software distribution requires its own TCP/IP stack with its own IP address and network ports.**

Solution:

- Deploying perfSONAR Testpoint in a docker container with **MacVLAN networking** solved this issue.
- Several IP addresses can be ssigned to the same network interface
- A MacVLAN network interface allows the Docker network to have its own NIC, MAC, IP address and TCP/IP stack.
- This network isolation method allows the MacVLAN interface to effectively use the bandwidth of the host NIC to achieve nearly lossless network performance (full 100 Gbps).

A national initiative of the Department of Science and Innovation and implemented by the CSIR
science & innovation
Department:
Science and Innovation
REPUBLIC OF SOUTH AFRICA
CSIR
Touching lives through innovation

SANReN Data Transfer Service

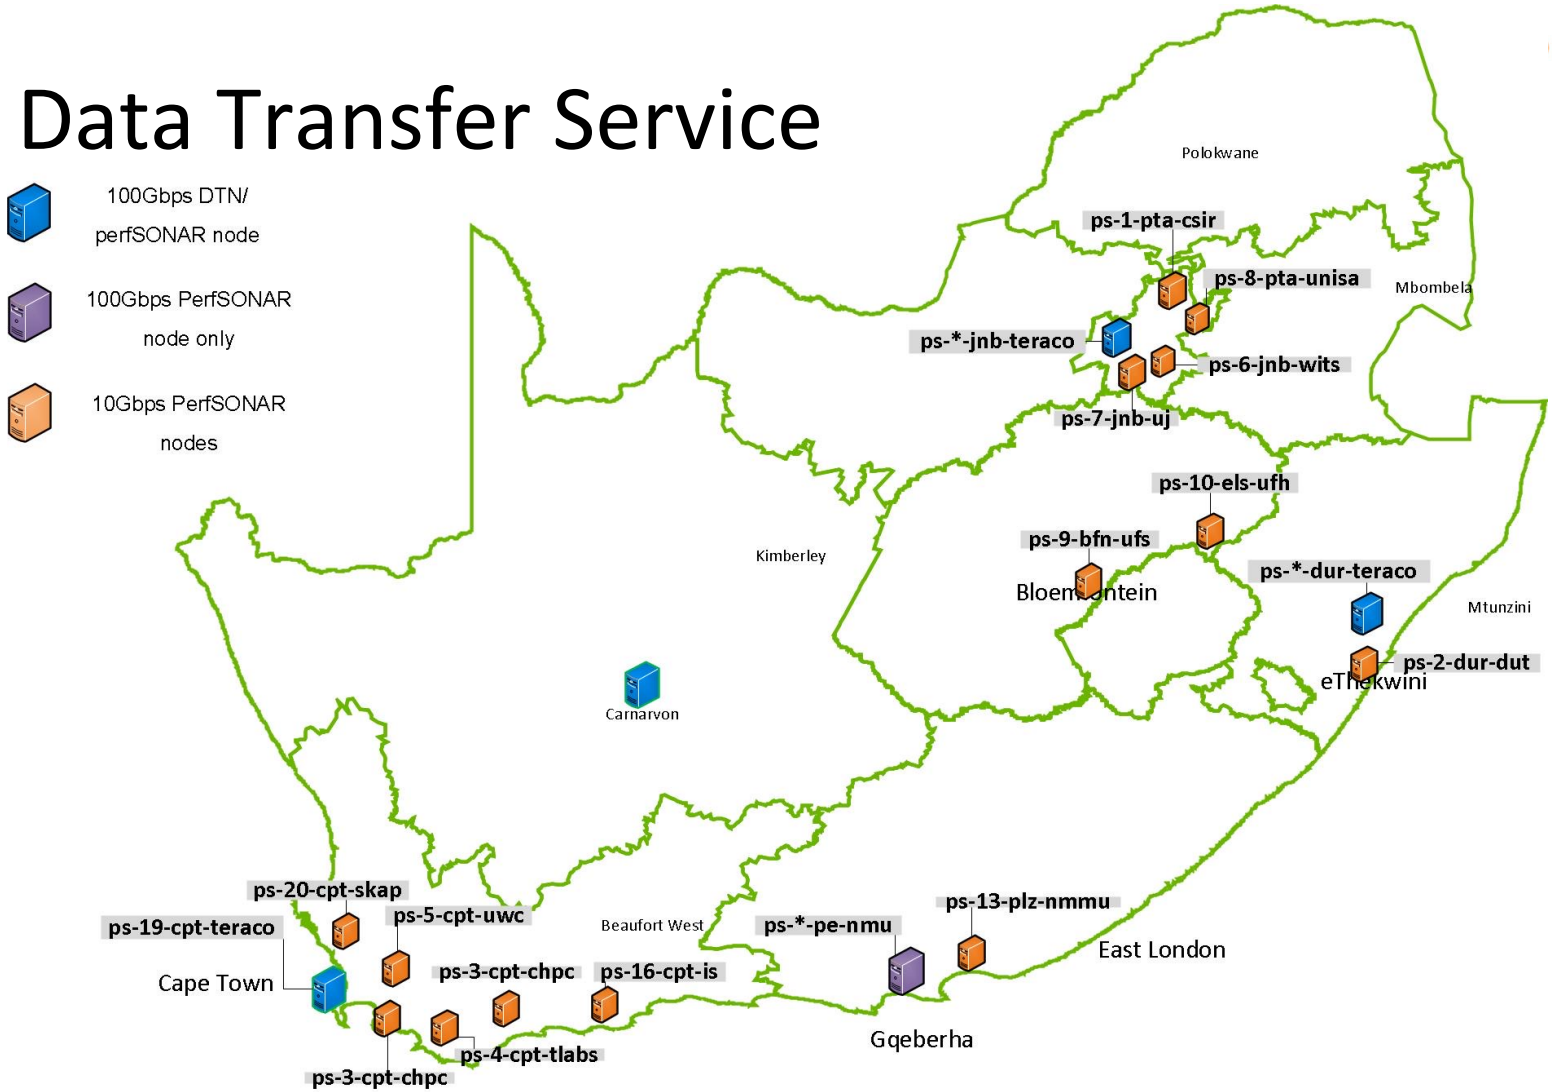

Benchmarking our 100Gbps DTNs

- We benchmarked our new 100Gbps DTNs against well performing 100Gbps DTNs in the US

- Thanks to Engagement and Performance Operations Center (EPOC) Data Transfer Testing/Data Mobility Exhibition (DME)

(supported by the National Science Foundation under Grant No. 1826994 info@epoc.global)

A national initiative of the Department of Science and Innovation and implemented by the CSIR

science & innovation
Department:
Science and Innovation
REPUBLIC OF SOUTH AFRICA

CSIR
Touching lives through innovation

Benchmark and test procedure results

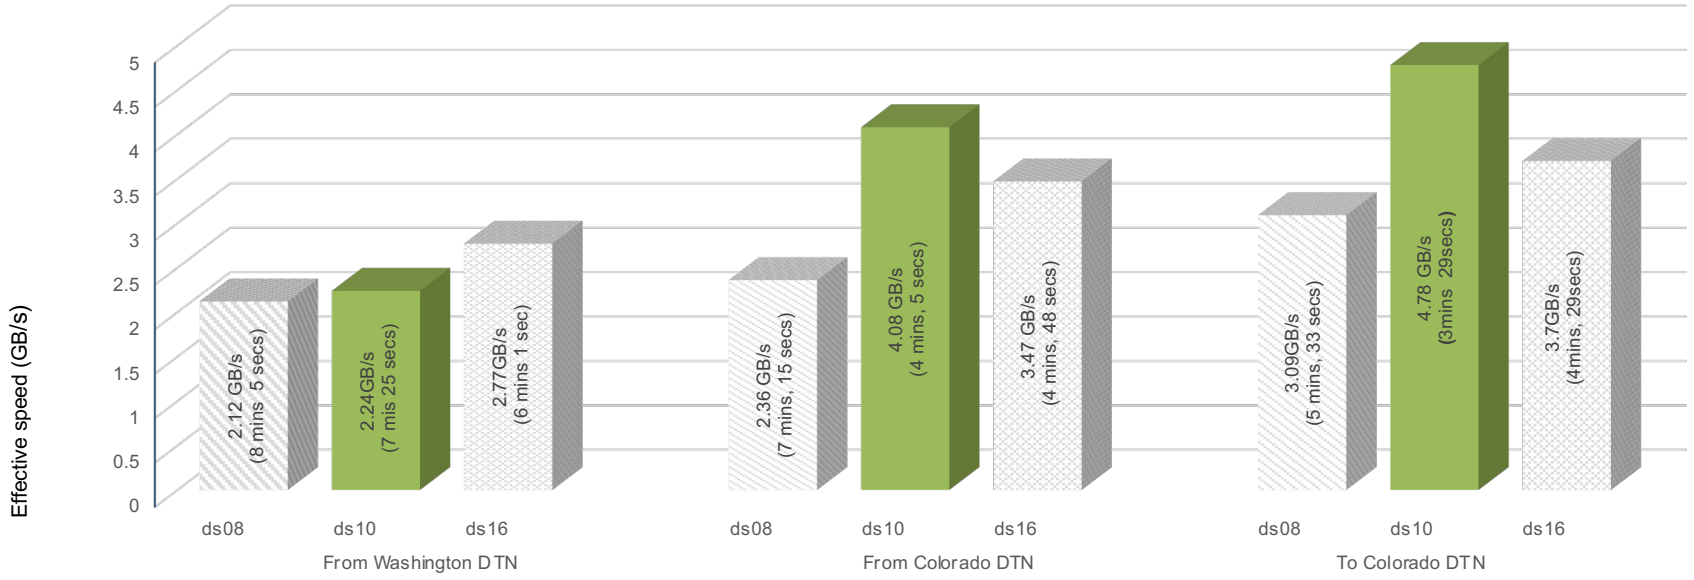

Results of Globus transfers from Cape Town DTN

A national initiative of the Department of Science and Innovation and implemented by the CSIR

science & innovation
Department:
Science and Innovation
REPUBLIC OF SOUTH AFRICA

CSIR
Touching lives through innovation

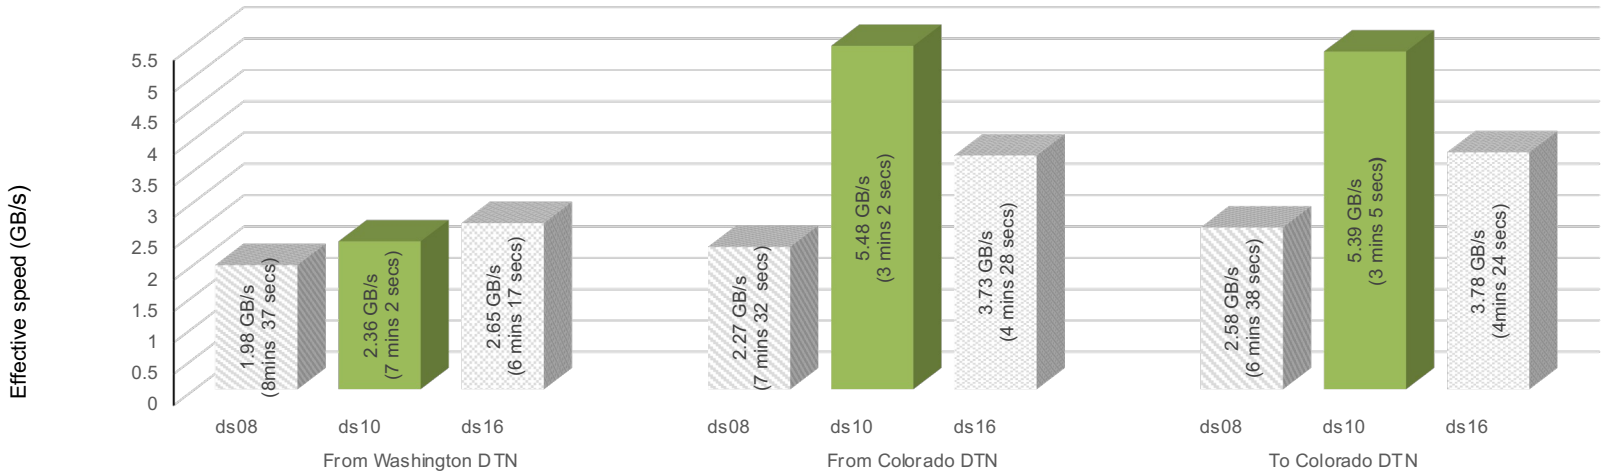

Results of Globus transfers from Johannesburg DTN

A national initiative of the Department of Science and Innovation and implemented by the CSIR

Benchmarking our 100Gbps DTNs

According to ESnet scorecard, an 'acceptable' result is 1TB in 1 hour. (https://fasterdata.es.net/DTN/data-transfer-scorecard/)

- Best test results to National Center for Atmospheric Research (NCAR) Globally Accessible Data Environment (GLADE)

- Cape Town DTN best test result is

  1TB data transfer from Colorado (NCAR GLADE) to Cape Town shows 3 mins, 29 seconds (4.78GB/s)

- Johannesburg DTN best test result is

  1TB data transfer from Colorado (NCAR GLADE) to Johannesburg shows 3mins, 2 seconds (5.48GB/s)

A national initiative of the Department of Science and Innovation and implemented by the CSIR

science & innovation
Department:
Science and Innovation
REPUBLIC OF SOUTH AFRICA

CSIR
Touching lives through innovation

Use case:

Physics data for the H-Line, Low Energy Nuclear Astrophysics Beamline project

Data size:	10TB
From:	Tandetron Facility, iThemba Labs, Cape Town.
To:	Texas A&M University, Texas, USA and INFN-LNS, Bologna, Italy.

- Network analysis and troubleshooting using the perfSONAR network toolkit
- Liased with US and Italian NRENs (ESnet and GARR) respectively.
- Liased with TENET to correct path routing and troubleshooting network links from iThemba Labs to Cape Town.
- Made available the SANReN Cape Town data transfer node for data sharing.
- Approximately 10TB of data was transferred from the experiment, over a week.

A national initiative of the Department of Science and Innovation and implemented by the CSIR

science & innovation
Department:
Science and Innovation
REPUBLIC OF SOUTH AFRICA

CSIR
Touching lives through innovation

Use case:

African Health Institute (AHRI)

Data size: 2TB
From: Africa Health Instituted (AHRI), Durban, South Africa
To: Harvard University, Research Computing HPC, USA

- Attempted to upload 2TB of genome FASTQ files from sftp.ahri.org to the Harvard server
- Initial current speeds that they were achieving (700kbps), it would have taken approx. 35days to complete the transfer.
- Using the SANReN 100Gbps Johannesburg DTN a peak transfer rate of 8Gbps was achieved and the transfer took approximately 40-45min to complete.

Data size: 1TB
From: Africa Health Instituted (AHRI), Durban, South Africa
To: Colorado State University, USA

Use case:

Wits University – Global Change Institute

Data size: 12TB

From: Australian National University National Computational Infrastructure (NCI)

To: Global Change Institute, University of Witwatersrand, South Africa

- The data was retrieved successfully in July via a SANReN DTN.

A national initiative of the Department of Science and Innovation and implemented by the CSIR

science & innovation
Department:
Science and Innovation
REPUBLIC OF SOUTH AFRICA

CSIR
Touching lives through innovation

Use case:

Very Long Base Line Interferometry (VLBI) experiments

Data size: 80TB for processing and then output 100GB.

From: Manchester, United Kingdom and University of Cape Town, South Africa

To: University of Pretoria (UP), South Africa

- The project is dedicated to uncovering supermassive black holes
- UP has a dedicated computing cluster for data processing 10s -100s of galaxies.
- Currently limited by the UP international bandwidth and computational cluster.
- These data transfers are conducted using a tool called JIVE for VLBI.
- Attempting to leverage SANReN's data transfer node infrastructure to replace/supplement the UP cluster.

A national initiative of the Department of Science and Innovation and implemented by the CSIR

science & innovation
Department:
Science and Innovation
REPUBLIC OF SOUTH AFRICA

CSIR
Touching lives through innovation

Transfers to/from CHPC in SA

- From the 100Gbps SANReN Data Transfer Nodes
- 1TB in various data compositions transfers in 15-26 mins.
- Effective data transfer speed of between 652.34MB/s – 1.12GB/s
- Sometime slower than international links – due to possible congestion on local links.

Supercomputing23 participation

- Flood the gates demo at Caltech booth
- The AmLight-SACS project established a 100G pathway using the South Atlantic Cable System (SACS) to South Africa (Cape Town) via Brazil and Angola

- We aimed to fill up link between South Africa and the United States to demonstrate the capacity of the 100Gbps AmLight link and SANReN's new 100Gbps data transfer nodes

- This was achieved through iperf testing from US to two new South African 100Gbps Data Transfer Nodes in Cape Town and Johannesburg Teraco data centres

A national initiative of the Department of Science and Innovation and implemented by the CSIR

science & innovation
Department:
Science and Innovation
REPUBLIC OF SOUTH AFRICA

CSIR
Touching lives through innovation

Supercomputing24 participation

1. Short paper accepted at INDIS (Monday 18th November, 12.18pm – 12.24pm)
2. Presentation at SciNet theatre (Tuesday 19th November, 16.40pm – 17.00pm)
3. Exhibit at the California Institute of Technology (Caltech) booth 845

   (Tuesday 19th November – Friday 22nd November)
4. Presentation at the California Institute of Technology (Caltech) booth 845

   (Date, time TBD)
5. "Flood the gates " demo, Caltech booth 845

   (Thursday 21st November, time TBD)
6. Manning HPC Around the World (Date and time TBD)

A national initiative of the Department of Science and Innovation and implemented by the CSIR

science & innovation
Department:
Science and Innovation
REPUBLIC OF SOUTH AFRICA

CSIR
Touching lives through innovation

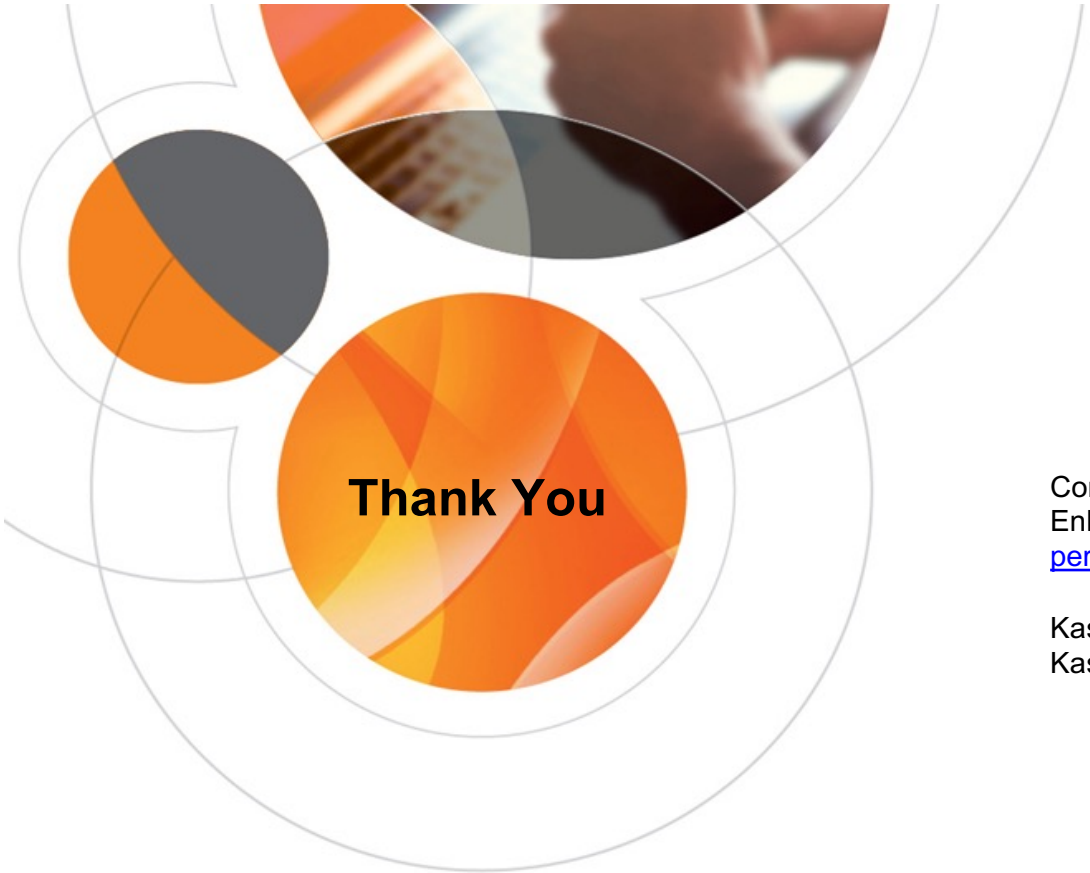

Thank You

Contact details: SANReN Performance Enhancement Response Team
pert@sanren.ac.za

Kasandra Pillay
Kasandra@sanren.ac.za

A national initiative of the Department of Science and Innovation and implemented by the CSIR

science & innovation
Department:
Science and Innovation
REPUBLIC OF SOUTH AFRICA

CSIR
Touching lives through innovation